

Information Age

Business & Technology Analysis

[Home](#) | [Events](#) | [Webinars](#) | [Blogs](#) | [Podcasts](#) | [Top 5](#) | [Contact us](#)


[Information Age > Cutting edge?](#)
[ARCHIVE](#) [HARDWARE](#)

Cutting edge?

'Blade' servers claim to squeeze as much computing power into as little floor space as possible. How do they work, and can they offer a scalable and manageable alternative to the way organisations currently buy server power?


19 March 2002

In his 11 years as chief technology officer (CTO) of investment bank Goldman Sachs, Vern Brownell developed a first-hand knowledge of how to manage a complex, distributed data centre, often within exacting budgetary constraints.

As technology became more central to Goldman Sachs' business, Brownell's IT department needed to move increasingly quickly to deploy new applications. As CTO, Brownell realised that the obstacles to achieving this were primarily physical: acquiring server hardware, installing it, and connecting it to networks and storage. Furthermore, in the face of a mounting skills shortage, finding the system administrators to perform these labour- and time-intensive tasks became increasingly difficult. Routinely, he says, it took six weeks or more to deploy a new machine.

"Over the years, I witnessed the pain points most CIOs suffer the most," says Brownell. But when he approached the companies that supplied Goldman Sachs' hardware to come up with a satisfactory solution to this challenge, they were not forthcoming. Brownell decided to put his own ideas to work: in 2000, he formed his own company, Egenera.

Egenera is one of a host of both new and established hardware companies that has latched on to the 'pain points' that organisations face in managing highly distributed and complex server architectures. As businesses grow and an increasing amount of corporate applications are accessed over the Internet, organisations typically add more servers. As a result, data centres are spilling over with



Welcome - Log in

Username:

Password:

[Go >](#)

☒ E-MAIL A FRIEND

☐ PRINTER FRIENDLY

☐ SAVE ARTICLE

☐ FOR REPRINTS

Blogs

Information

Enterprise Software

Unwrapped

Information Age 1c

Information Age

Current issue

Subscribe to print edition

About Information Age

Register

For full site access

For blogs and news

feeds

Briefing rooms

Business Applications

Business Intelligence

Communications

Data Centre


Data Integration

Information Management

IT Infrastructure

IT Services

Security and Continuity
SOA and web services
Storage



Ads by Google

Blade Servers
Blade IBM
Blade HP
Blade Dell
Blade into Servers

servers that have become difficult to manage and costly to maintain.

The solution Egnera is proposing – alongside industry heavyweights such as IBM, Compaq, Hewlett-Packard and Sun Microsystems – is the blade server.

The basic premise of blade servers, which contain the equivalent of an entire server on a single board, is to squeeze as much computing power into as little space as possible (see box, What is a blade server?). By installing these servers in their organisations, claim the blade vendors, IT departments can slash the total cost of ownership (TCO) of their server architecture, significantly reduce application deployment times, and more easily manage distributed servers.



Vern Brownell,
Egnera:

"Ultimately it will
be the way that
all servers are
built."

"The blade server form factor has been serving telecommunications companies for years – it's not new," explains Tom Bradicich, director of architecture and design for IBM's eServer X-Series division. "What is new is that the technology is moving into the mainstream enterprise server market."

Currently, IT departments tend to add more server power in 'racks', deploying one- or two-processor servers as Internet demand increases. While these rack systems are relatively dense compared with high-end, multi-processor servers, says Bradicich, they are difficult to manage, can take hours or even days to deploy, and if one goes down, it is hard to diagnose precisely where the problem lies because they are so dispersed.

But how do blade servers differ from traditional servers, other than cosmetically? The key benefit of a blade server, argues Bradicich, is ease of management. "You can quickly replace a blade and put it back into the rack like a book in a bookshelf. This means organisations can upgrade and replace applications in minutes instead of minutes."


Because they are supported by sophisticated management software – typically contained on one of the blades themselves and linked to a central administration console – managing application performance, load balancing and allocation of resources is all centralised.

For organisations that have highly dispersed computing architectures, this could allow them to make significant cost savings in terms of both staffing and property expenses. One UK financial institution based in the City of London, for example, plans to replace its existing data centre – consisting of some 56 computer rooms dispersed at various points around the City, with just two data centres located 20 miles apart, by centralising its architecture onto blades.

Furthermore, the fact that a greater number of servers can be contained in the same chassis, sharing cooling, power and cabling resources, means the cost and provision of these basic resources can be amortised over a greater number of systems. With traditional servers, organisations pay for these features on a server-by-server basis.

Despite their apparent benefits, however, blade servers are still very much in the 'early adopter' phase. Although a host of start-ups have launched blade designs during the past several quarters, only two of the major hardware vendors – Hewlett-Packard (HP) and Compaq – have actually brought blade products to market so far, although IBM, Dell and Sun Microsystems are expected to follow during the second half of 2002. According to conservative estimates from analyst company Gartner, worldwide shipments of blade servers will reach just 900,000 units in 2002 and will only surpass 1 million by around 2006.

One of the key reasons for this slow adoption is doubts over how these servers will interoperate with users' existing infrastructures. "The underlying hardware and application software will follow the same standards. But the designs of the blades themselves will differ from in the past," explains Gartner



analyst Adrian O'Connell. Because the low-end server market has become so commoditised, he adds, vendors may find it in their best interests not to promote interoperability between their hardware or management software and those of other vendors in order to achieve a higher degree of customer lock-in.

"This is something users have to weigh up for themselves," advises O'Connell. "They need to ask themselves, 'Six months down the line, am I constrained to stay with this vendor?'"

So how should organisations approach integrating the new blade architecture into what they have already? Andreas Knoepfli, vice president of the industry standard server group for Compaq in Europe, advises prospective users to ensure blades comply with existing industry standards in areas such as systems management software, IP connectors and interfaces with storage systems.

"Some of the start-ups in this market are suggesting that companies rip out their existing infrastructure. [Blade servers] should not replace what they already have, [they] should be incremental to what is there already," he advises.

Organisations should also bear in mind what type of applications they intend to run on the blade servers. As Peter Roberts, UK sales director of distributed computing specialist Platform Computing, points out, the fact that the vast majority of blades run either Windows or the open source operating system Linux means that organisations will be restricted to running Linux- and Windows-compatible applications.

Both HP and Compaq concede that they anticipate the greatest adoption of their blade products to be at the front end of the data centre, performing simple web caching, load balancing and web serving functions. "We see blades as a more cost-effective way to deploy thin rack servers than the way they are deployed today. We're not trying to replicate our high-end machine," says Jon Jacob, European mid-range server marketing manager at HP.

Although most of the early models of blade servers operate on a single processor, the introduction of two- and four-way blades over the next few quarters will allow organisations to scale up the kind of applications they run on them. Dell, for example, plans to launch its first blade server product in the second half of 2002, which will be dual-processor. "You could have a full office on one single chassis, containing a virtual private network, a groupware application such as Notes or Exchange, plus all the usual server functions," explains John Bailey, UK server business development manager at Dell.

Analysts at Gartner believe that the introduction of blades is a "market disrupter", a "bump" in the server market cycle that will allow users to take advantage of new technology and vendors to take share from competitors. But with this bump in the market comes a need for caution. "The demand for front-end web servers has been critical in driving the need [for blade servers], but in reality, the market is only just beginning to understand how blade servers might be deployed and used," says O'Connell.

Brownell of Egenera has, naturally, more conviction. Having given up his CTO job at Goldman Sachs to develop a new server architecture that he hopes will ease management and lower costs for IT departments, he has much to lose if the market does not take off.

"Blade is definitely a step in the right direction," he enthuses. "Ultimately, it will be the way servers are built."



Andreas Knoepfli, Compaq: "Blade servers should be incremental to what is already there."

Although the main attractions of the blade server may seem simply cosmetic – organisations can reduce the amount of data centre space taken up by servers, cables and networking equipment – as they gain acceptance analysts believe they will play an important role in a couple of wider computing trends.

One of these key trends is self-healing, or 'autonomic' computing, where computer systems automatically diagnose and resolve problems in much the same way as the human body reacts to pain or cold conditions.

As the management software and intelligent allocation of computing resources behind blade servers mature, this will mean much of the diagnostic work usually performed by systems administrators will be performed automatically.

A second key computing trend in which blade servers are likely to figure is 'grid' computing. This is where huge numbers of computers dynamically share processing resources across private networks, and in some early instances, the Internet.

Much of this work – often supporting processor-hungry applications such as scientific research – is currently done across high-end Unix, mainframe or even supercomputer systems. As these organisations become more constrained for space, however, it is conceivable that they could run these grid-optimised applications across large clusters of blades.

The Sanger Institute in Cambridge, UK, for example, which has been using an early grid computing architecture to work on the Human Genome Project, has recently received funding to double the amount of server power it is dedicating to its research. If it does this based on its current architecture, however, it will have to build a new data centre. With blades, it would not have to waste research money on expanding its premises.

Peter Roberts, UK sales director of Platform Computing, the company that has supplied the distributed computer management infrastructure to support Sanger's Human Genome Project research, says this is a prime example of where blade servers would work well.

"There will always be applications that need to work on supercomputers, but the drive for reduced footprint and power consumption will force many organisations to write grid-enabled applications that work across huge clusters of blades, allowing them easier management as well as reduced costs," he says.

Like similar organisations of its size, US investment bank Credit Suisse First Boston (CSFB) has experienced a steady increase in the number of servers it requires to host its applications and, consequently, the number of administrators required to manage these systems. As its infrastructure grew, CSFB faced the challenge of ensuring its data centre infrastructure could support its business-critical applications, yet still keep operational costs down.

"With both data centre space and qualified personnel expensive and scarce, the only answer was server consolidation," explains Evan Bauer, chief technology officer for investment banking IT at CSFB.

In 2001, Bauer and his team began evaluating a number of platforms, including the option of simply extending the number of RISC-based Unix systems it had already. However, after performing a total cost of ownership analysis on how much it would cost to do this, Bauer found that if CSFB migrated its blade servers from start-up server vendor Egenera, which claims to offer comparable performance to a high-end Unix system on a thin blade server, the cost would work out at around a quarter of that of the equivalent RISC-based machines.

"A CSFB application that would have required 20 of our standard RISC Unix servers can be deployed on just four Egenera blades," says Bauer. Furthermore, because the blade chassis itself contains all the necessary cabling, network switches and routers, it cuts down on the amount of data centre space needed.

Initially, Bauer introduced the blade architecture to support CSFB's global order-routing system, which processes an average of 20 million transactions per day.

"With our first application, we've saved over \$900,000 in hardware and software alone. As we deploy multiple applications across the architecture, we plan to save millions of dollars," says Bauer.

A number of additional applications are slated for roll-out during 2002. One of the future roles of the Egenera systems, for example, will be as a disaster recovery resource. The company plans to directly link the Egenera blades to remote back-up software and storage systems from EMC across a number of the company's regional data centres.

This has led some industry analysts to term the blade server phenomenon 'ultra-dense computing'. Unlike their rack server relatives, blade servers are comprehensive computing systems – typically based on Windows or Linux – that include processor, memory, network connections and all the associated electronics on a single motherboard. Networking, storage and systems management blades sit in an enclosure alongside blades serving applications or web pages. Because so many blades can be hosted in a single rack, they can share common resources such as power supplies and cooling fans, cutting down on operational costs. There is also very little cabling in comparison with traditional servers.

The unique selling point of blades, claim vendors, is their ease of management and flexibility. Using sophisticated management software, systems administrators can install and configure specific blades to serve particular applications remotely and in a matter of minutes. They can also take servers off certain applications if they require less processing power, and add power at times of peak demand.

[Back to main text](#)

By Jo Faragher, faragher@infoconomy.com

[Back to top](#) ▲

[SITE MAP](#) | [TERMS AND CONDITIONS](#) | [PRIVACY POLICY](#) | [ADVERTISING INFO](#) | [CONTACT US](#) | [ABOUT](#) | [ARTICLE DIRECTORY](#) | [USEFUL LINKS](#)
Copyright © Infoconomy 2006 | POWERED BY MySource